DUNE CCE Priorities Year2+

Michael Kirby, Fermilab/Scientific Computing Division

Nov 6, 2020

ASCR CCE All Hands Meeting

DUNE Computing Challenges

- FD produces 30 PB raw data per year
  - FD FE/DAQ has much greater bandwidth
- Time-extended trigger records present unique situation
  - 100s instead of localized event of 5.4 ms
  - after zero-suppression estimated to be 184 TB
- Processing challenge for Far Detector
  - initial 2D Deconvolution of spatially independent Anode-Plane Assembly
  - Hit Finding and ROI reduce trigger record from 6 GB
  - Extensive event reconstruction algorithm (3D matching, vertexing, calorimetry, particle identification)
  - Machine Learning Training & Inference
- Simulation - computational challenge of large “open” detector, plus all that ↑
- Prototyping two FD module designs
  - ProtoDUNE-SP - beam (2018) and cosmic-ray operations
  - ProtoDUNE-DP - cosmic-ray operations
  - more than 4 PB of raw data for SP and DP

Far Detector SP Module

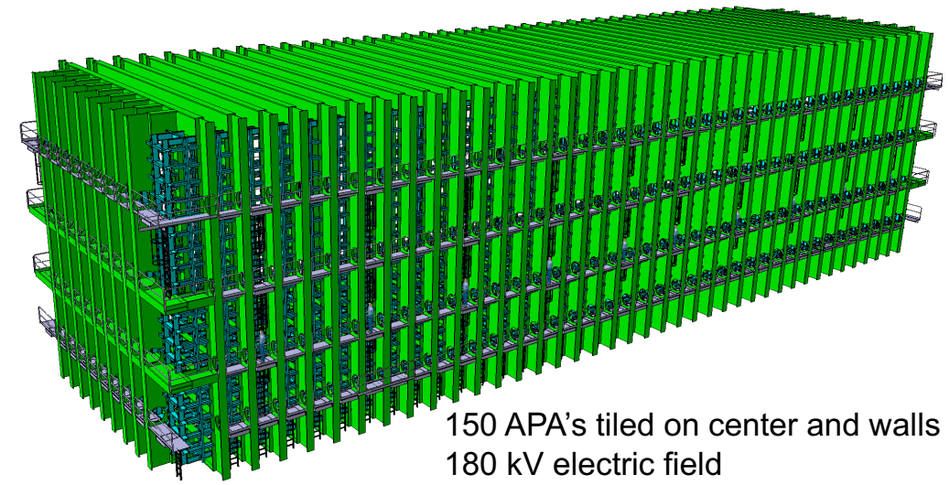

150 APA’s tiled on center and walls
180 kV electric field
17 kT of Argon

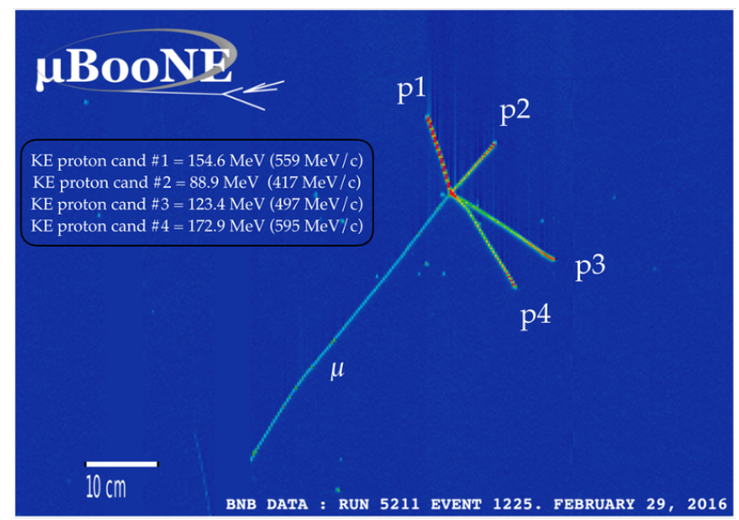

Near Detector Physics

- Three separate detectors with diverse readout structures
  - Requires triggering and event building across detectors
- Designed to address:
  - oscillation neutrino flux measurement
  - Neutrino-nucleus cross section measurements
  - Beam/Target monitoring
- information density much higher in ND raw data
  - derived dataset volume
  - large simulation statistics
  - multiple detector configurations

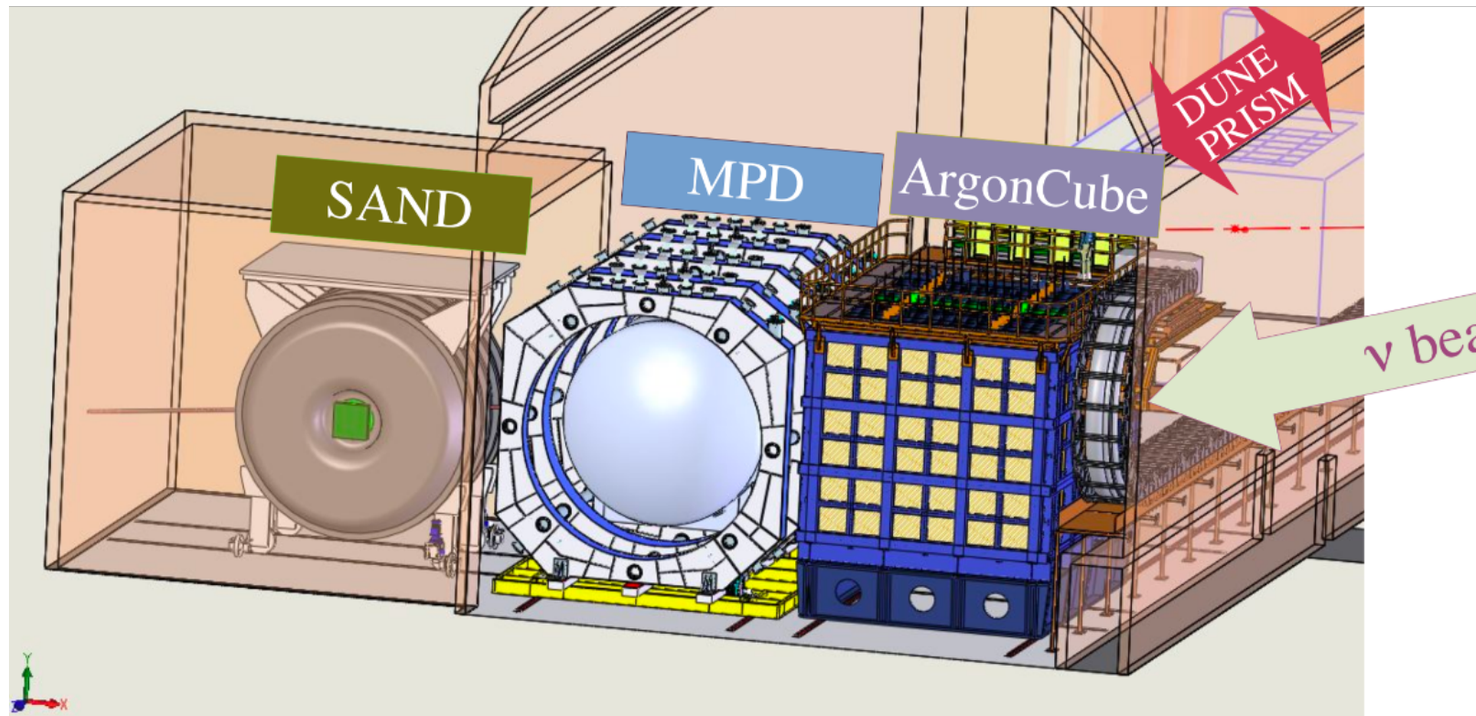

DUNE Priorities going forward for CCE PPS

- Continued progress on WireCell studies towards the goal of first metrics within Kokkos
  - full suite of modules from signal processing, hit finding, tracking, vertexing, etc. provides end-to-end reconstruction and complete workflow
  - potential to study individual modules to have better understanding of portability packages capability for types of algorithms
  - FNAL priority to find additional focused effort for this project to help complete integration and initial metric evaluations
  - WireCell team considering new design of data structures with focus on multithreading/parallelization/accelerators
- Considering a secondary workflow which is the portability of CUDA based Machine Learning algorithms
  - Research into Image Processing ML for LAr TPCs is very active and builds upon industry algorithms and advancements
  - in contact with collaborators that have CUDA-based modules that might be more comparable with Patatrack or FastCaloSim CUDA implementations

WireCell

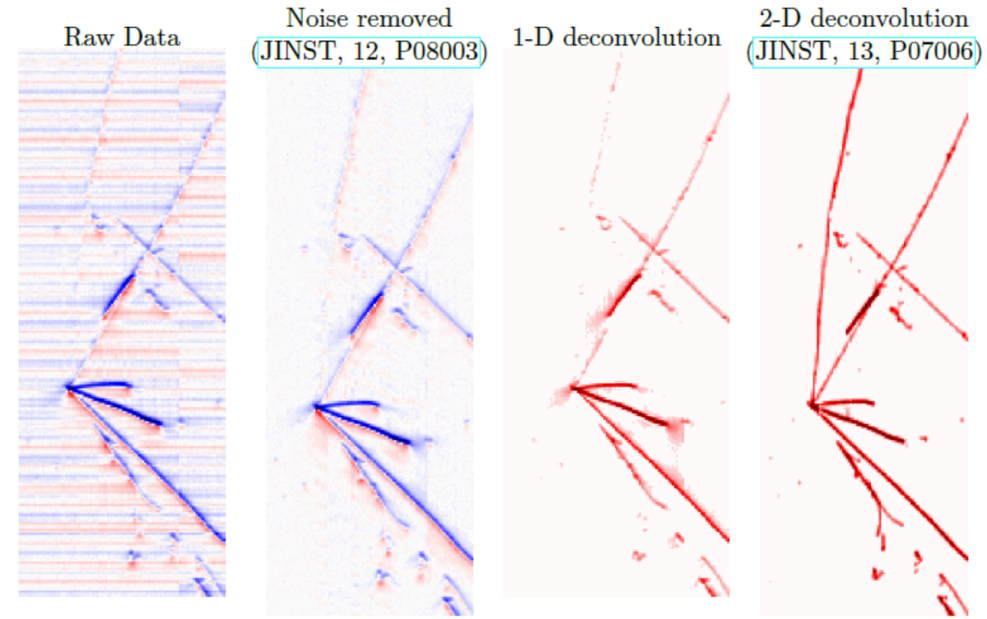

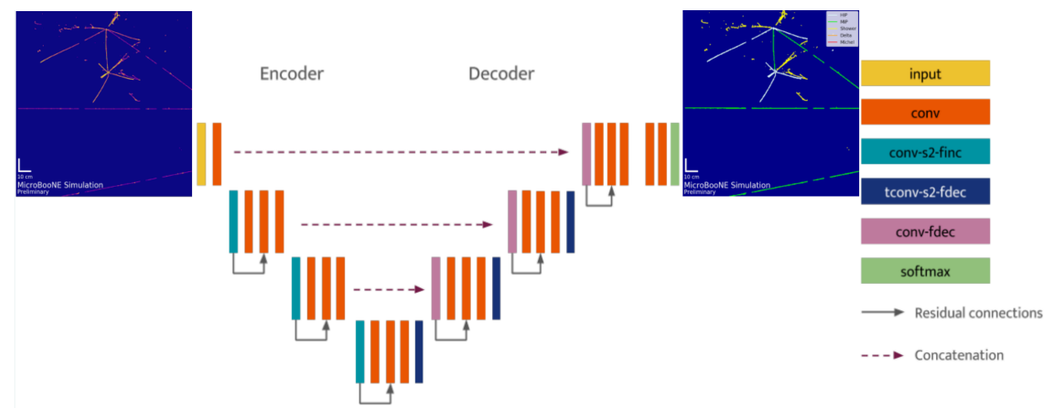

SparseSSNet - Ran Itay (SLAC)

DUNE Priorities for CCE IOS

- DUNE is actively investigating using HDF5 for raw data format from DAQ
  - recognize the already demonstrated performance from SciDAC for HPC
  - trying to understand the performance, advantages, and limitations within DAQ, HTC, and interactive processing
  - potentially utilize during ProtoDUNE II operation in second half of 2022
- Understanding of the current workflow I/O pattern has significant implications for computing model
- Recent results in GPUaaS shows real promise

Inference with GPU as a Service

- arXiv:2009.04509
- CNN EmTrackMichelId module in ProtoDUNE SP reconstruction
- utilized SONIC approach to have Fermigrid CPU jobs offload inference to Google Cloud GPUs deployed with Kubernetes
- CPU only - 7 s preprocessing, 213 s/evt running 55k inferences
- When using GPUaaS 13 s/evt
  - GPU ~30k inferences/s -> 1.7 s/evt
  - 7s preprocessing, 4.5 s I/O

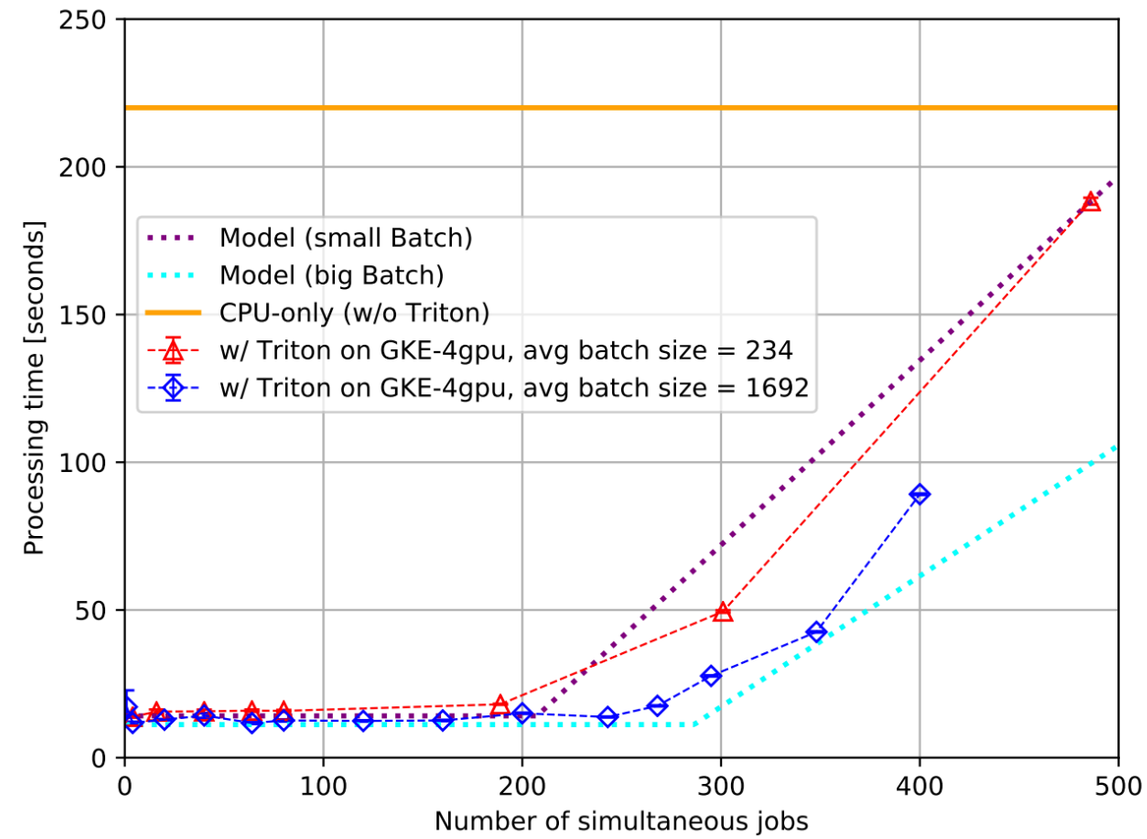

Complex Workflows

- raw data processing may happen in large, weekly batches on HPC from HDF5 stores
- feed into HTC based event reconstruction
- utilize GPUaaS for inference
- track all of it in flight
- supernovae processing might go straight to HPC and depend on availability

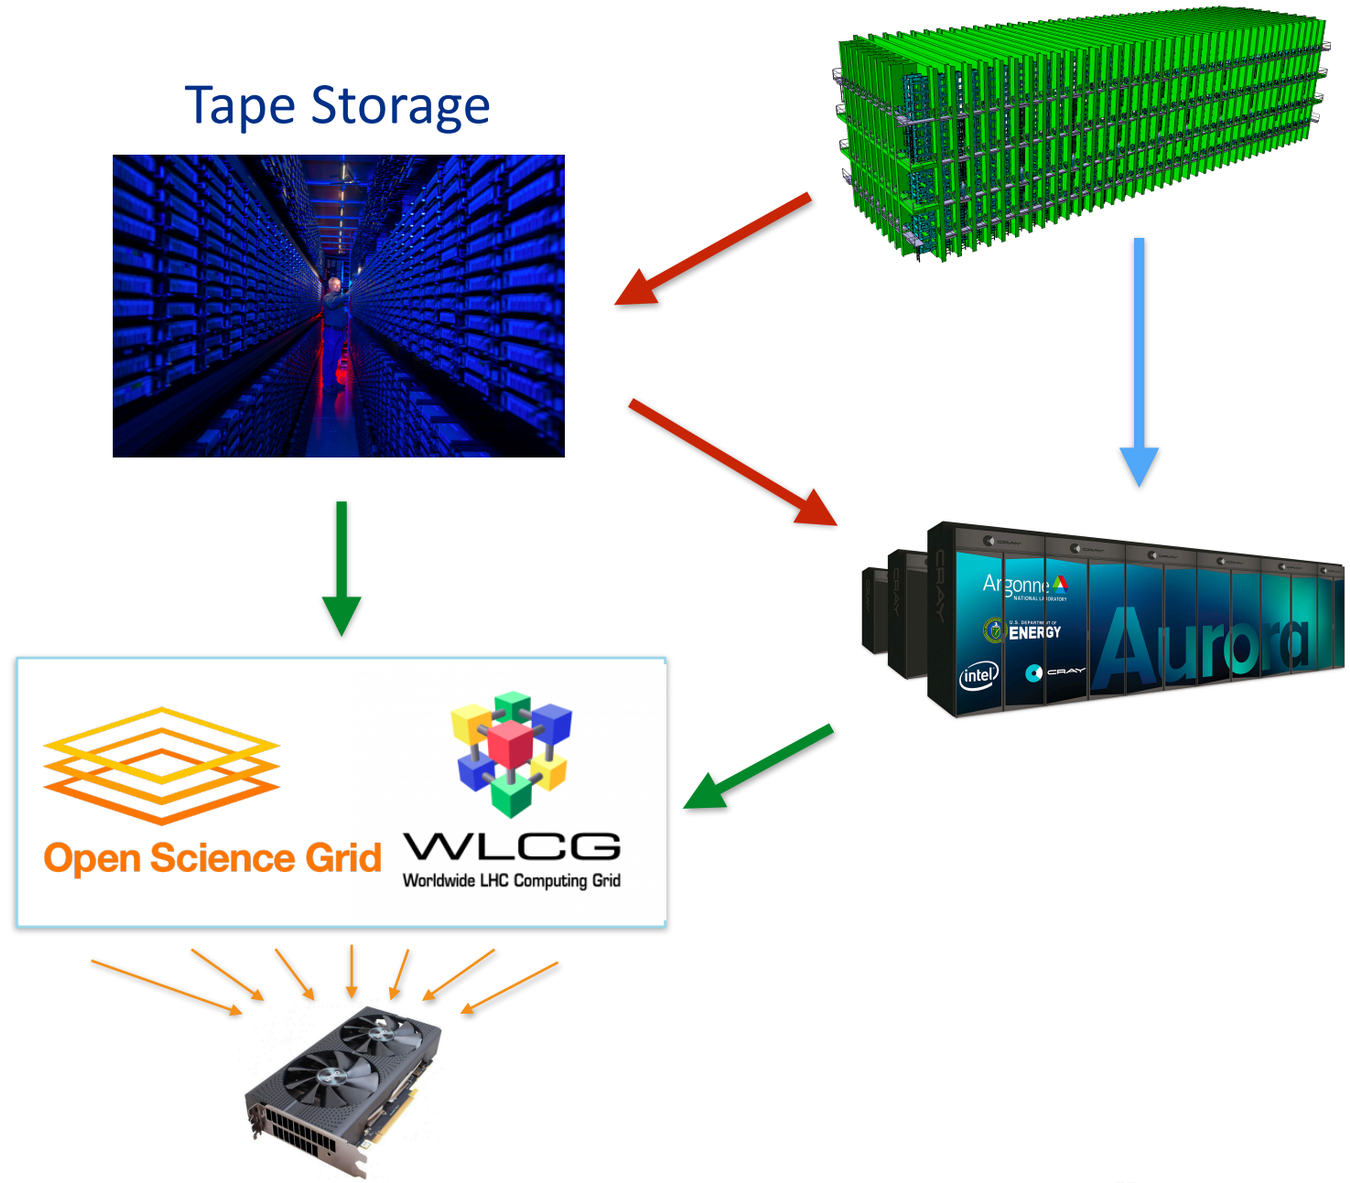

DUNE Overall Priorities

- Detectors will not exist for another 7 years
- We haven’t made all the choices that will define the choices we will need to make
- CCE - PPS - important for us to learn about how to make choices about portability for both production and analysis algorithms
- CCE - IOS - HDF5 as potential raw data format for processing on HPC
- Complex workflows for localized-readout raw data, and for supernova 200 TB triggers